

An Empirical Study of Pricing Strategies in an Online Market with High-Frequency Price Information

Sara Fisher Ellison
M.I.T.

Christopher M. Snyder
Dartmouth College

May 2011

Abstract: We study competition among a score of firms participating in an online market for a commodity-type memory module. Firms were able to adjust prices continuously and prices determined how the firms were ranked and listed (lowest price listed first), with better ranks contributing to firms' sales. Using a year's worth of hourly data, we document the pricing dynamics, cycles, and other patterns in this market. We then characterize empirically the factors which drive price changes, noting clear evidence of firm heterogeneity in the choice of pricing strategy. Finally, we develop a framework for simulating counterfactual market settings, using the simulations to examine counterfactuals involving different mixes of firms according to pricing strategies.

JEL Codes: L11, C73, D21, L81

Contact Information: Ellison: Department of Economics, M.I.T., 50 Memorial Drive, Cambridge, MA 02142; tel. (617) 253-3821; fax. (617) 253-1330; email sellison@mit.edu. Snyder: Department of Economics, Dartmouth College, 301 Rockefeller Hall, Hanover, NH 03755; tel. (603) 646-0642, fax. (603) 646-2122, email chris.snyder@dartmouth.edu.

Acknowledgments: The authors are grateful to Hongkai Zhang for superb research assistance, to Glenn Ellison for a number of useful conversations, and to Florian Englmaier, Georg Gebhardt, Chakravarthi Narasimhan, Mo Xiao, and various seminar participants for useful comments.

1. Introduction

There is a vast and rich theoretical literature on the dynamics of repeated games. One conclusion of this literature is that many different market outcomes are possible. Full collusion, collusion on a focal price, cyclical markup patterns, and prices exactly tracking costs, are among the possibilities. This lack of a clear theoretical prediction suggests an important role for empirical analysis. In any particular setting, theory might suggest that certain outcomes are more or less likely, but the empirical documentation will ultimately be important. We are interested in the market outcomes in one particular empirical setting, an online market with a large number of firms, ranked by price, with highly visible and easy-to-change prices. We want to understand the effect that this transparency and immediacy has on the firms' abilities to maintain markups, and, more generally, the factors that affect firms' decisions to change prices. We also want to determine whether the particulars of our empirical setting can give rise to cyclical patterns, and, if so, of what variety.

The empirical literature on repeated games has often been hampered by the paucity of complete, high-frequency pricing data over long periods of time. Notable exceptions include work on the Joint Executive Committee railroad cartel by Porter (1983) and Ellison (1994), the Sugar Institute sugar cartel by Genesove and Mullin (2001), and British shipping cartels by Scott Morton (1997); research on Edgeworth cycles in retail gasoline pricing (Eckert and West 2004 and Noel 2007b among others), and more recent work on Edgeworth cycles in online advertising auctions by Edelman and Ostrovsky (2010) and Zhang (2010).¹ It is interesting to note that laws against collusion and other types of price fixing may have significantly hampered the study of pricing dynamics by academics because they have made the gathering of detailed pricing data by rivals a *de facto* suspicious activity. Regulated

¹The data on the Joint Executive Committee, the Sugar Institute, and the British shipping cartels were gathered by the market participants in a pre-Sherman Act period where, because collusion was legal, gathering detailed pricing data on rivals would not have been viewed as suspicious. The retail gasoline pricing data used by Noel were painstakingly gathered by hand twice every day on a commute to and from work. The data on online advertising auctions were gathered by Yahoo! and Google, who were actually running the auctions.

industries or industries in the process of deregulation, such as the wholesale electricity market, often gather detailed pricing data, but these data may be less useful for the study of unfettered competition. Finally, there are many markets that might be ill-suited to the empirical study of repeated pricing games simply because of their institutional structure. Markets where prices are hard to change or hard to observe may yield less interesting pricing dynamics, for instance.

Web-scraping technologies, of course, have started to enhance the availability of pricing data. They have made detailed pricing data fairly easy to collect, and researchers have started to exploit them. Baye, Morgan, and Scholten (2004) is an excellent example of a study which used detailed pricing data collected on the web over a period of time, although their focus was not on the firms' repeated game strategies. We, too, have gathered most of our data by scraping firm and pricing data from a price search website for computer components. Furthermore, institutionally, we have a setting amenable to the empirical study of repeated games. In particular, we have identified 47 firms fairly active in selling a particular commodity-type memory module over the internet during one particular year. They list their prices for the item on a price search website and can change their listed prices at will. Prices are observable by all market participants. Wholesale costs are volatile, so firms must change their prices sometimes or they will either be selling at large losses or get bumped off the list of the twenty-four lowest prices. Consumers focus mostly on price-based rank of the firms when making their purchase decisions,² so firms are acutely aware of their rivals' prices. The transparency and immediacy of this empirical setting, along with the relative homogeneity of the products offered, have the potential to result in interesting pricing dynamics.

Our goal with this paper is three-fold. First, we want to document the pricing dynamics, cycles, and other patterns in this market. Second, we want to characterize empirically the factors which drive price changes, noting clear evidence of firm heterogeneity. Finally,

²See Ellison and Ellison (2009a) for demand estimates in this market.

we would like to develop a framework for simulating counterfactual market settings. In particular, we will examine counterfactuals arising from different mixes of the types of firms participating in the market, but we note that this framework is sufficiently flexible to allow for a large variety of types of counterfactuals.

One of the striking aspects of this market that we note is firm heterogeneity. Standard economic reasoning would suggest that our setting would engender homogeneity, not the heterogeneity we observe. For instance, the specific product on which we focus is fairly undifferentiated. We exploit the fact that the price search engine that mediates this market has predefined product categories, and we focus on one high-volume category, 128MB PC100 memory modules. The firms with the twenty-four lowest prices in this category are all selling generic, untested modules with unfavorable warranty and return terms. Furthermore, brand names of firms are relatively unimportant. Most consumers would have never encountered the names of any of these firms before. From talking to one of the business owners in the market, we gather that the competing firms are all small, share wholesale sources, and have similar sets of products. And yet we observe significant heterogeneity in the repeated game strategies that they seem to employ. We find this heterogeneity interesting enough to document, but we must remain largely agnostic about its origin. Even though the firms share wholesale sources, they might have different cost structures internally and may, therefore, face different marginal costs of fulfilling orders. They could enjoy different levels of managerial competence, which could account for the different strategies we observe. Finally, even if the firms were identical in all of these respects, it is possible that the heterogeneity we observe is simply the result of experimentation with different strategies or mixing. In fact, the business owner to whom we spoke acknowledged experimenting constantly with different website designs in addition to pricing strategies.³

Our paper proceeds in eight sections. Section 2 reviews the relevant literature. Section 3 offers additional details about the empirical setting. Section 4 describes the data set we use.

³In fact, he opened multiple websites selling identical products to facilitate this experimentation.

Section 5 discusses how we classify firms into different “strategic clusters” to account for the heterogeneity we observe. Section 6 presents and interprets our empirical model of firm pricing. Section 7 presents our simulations of counterfactuals of the mix of different strategic clusters. We conclude in Section 8.

2. Literature Review

The first part of this section provides a review of relevant theoretical papers in order to provide some background for our empirical work. The second part surveys the related empirical literature, highlighting our paper’s contribution.

2.1. Theoretical Literature

The theoretical literature provides a mixed view on the ability of firms to sustain high margins in our market. The repeated-game literature (see e.g., Fudenberg and Maskin 1986) suggests that the monopoly outcome is sustainable if firms have a discount factor close enough to 1. In our market, firms could adjust price instantaneously, and these price changes are immediately registered to rivals along with customers, allowing rivals to respond to price cuts with their own. Translated back into a discrete-time setting, the period lengths in our market effectively are arbitrarily short and the discount factor arbitrarily close to 1, suggesting according to the Folk Theorem of Fudenberg and Maskin (1986) that full collusion is a possibility. Other factors that textbooks point to as facilitating collusion are also present in our market, including a homogeneous product and firms without glaring asymmetries among them (see Tirole 1988, Section 6.1, for a discussion).

On the other hand, other considerations may suggest that collusion is difficult, providing a rationale for the more competitive prices that we will end up observing. First, Fudenberg and Maskin’s (1986) Folk Theorem indicates that collusion on the monopoly outcome may be a possibility but not a necessity, with a whole range of other outcomes possible, including marginal-cost pricing. How firms arrive at their repeated-game strategies and the market

arrives at an equilibrium in the presence of this multiplicity of equilibria has been the subject of active research inspired by Axelrod's (1984) seminal work. A second factor hindering collusion is the relatively large number of firms in our market: hundreds of firms were listed on the price-comparison website (our data covers a smaller group of the 24 lowest-price firms each hour, for a total of 47 with significant presence during the sample period). For many reasons, collusion is harder to sustain with more firms.⁴ Third, although in principle firms were able to adjust price instantaneously, in practice prices were not continuously changed, with lags of days or even weeks between price changes. Why firms showed inertia in price changes is in part the subject of our analysis. Davis and Hamilton (2004) offer four suggestions: menu costs, information-processing lags, customer acceptance, and strategic recognition of rival responses. To this list we would add a fifth: costs of continually monitoring market conditions, and computing best responses to these, leading to periods of inattention.

Interesting price patterns can emerge with staggered price setting in repeated games. The seminal paper in this area, Maskin and Tirole (1988), showed that if two firms set prices in alternating periods, there exist Markov-perfect equilibria exhibiting a focal price and also others exhibiting cycling prices (called Edgeworth cycles after Edgeworth 1925). Although we will not see marked Edgeworth cycles in our price series, firms' rankings in the price list will exhibit cycles for some firms reminiscent of Edgeworth cycles, and one of the goals of the analysis will be to characterize these rank cycles. Recent theoretical research has pointed to the robustness of the cycling equilibria. Wallner (1999) shows that Markov-perfect equilibria with price cycles can emerge even if the repeated game has a finite horizon. Eckert (2003) shows that while equilibria with cycling prices exist for a large range of firm sizes, focal-price equilibria require the two firms to be close enough in size. Work showing that asynchronous pricing will emerge in equilibrium includes Fishman (1992), Cahuc and Kempf (1999), and Lau (2001). Bhaskar and Vega-Redondo (2002) justify the restriction to Markov strategies,

⁴Tirole (1988) provides a general discussion. Chowdhury (2008) shows that the unique equilibrium of a repeated game with convex costs (related to Edgeworth's capacity constraints) in which firms can choose price and quantity is a grid point above marginal cost if the number of firms is large enough.

showing that all Nash equilibria must be in Markov strategies if agents have limited memory and strategies conditional on lengthening histories involve complexity costs. This paper is also related to the theme of limits to the reasoning and processing power of the price-setting agents mentioned in the previous paragraph as possibly leading to pricing inertia.

The price cycles that emerge in the literature cited in the previous paragraph are solely a function of firm strategy and not driven by fluctuations in market conditions. Of course, price cycles and other interesting movements may be generated by underlying fluctuations in demand and/or cost. Indeed these market fluctuations may affect the nature of firms' strategies, typically providing another factor hindering collusion (Tirole 1988). Classic papers on supergames in the presence of market fluctuations include Rotemberg and Saloner (1986), analyzing the case of observable shocks, and Green and Porter (1984), analyzing the case of unobservable shocks. Several more recent papers have added market fluctuations to repeated games with staggered pricing (Fishman 1992, Eckert 2004, Leufkens and Peeters 2008).

Much of the theoretical literature tends to focus on symmetric strategies, in contrast to our findings that firms adopt heterogeneous strategies that fall into identifiable clusters. The theory does suggest some rationales for heterogeneous strategies. Not knowing the strategies rivals are playing to say nothing of the best response to these strategies, different firms may be experimenting with different strategies, one or some of which may begin to emerge as dominant through survivorship or other evolutionary forces, along the lines of Axelrod (1984). In the model of Hansen *et al.* (1996), with duopolists in a Bertrand game with differentiated products, an unanticipated negative demand shock may elicit asymmetric responses as a decrease in one firm's price reduces the benefit to the other from changing its price, and it may decline to do so if price changes involve a fixed cost. In Lal and Matutes' (1989) model, manufacturers of a product bundle charge different prices for the individual components although the overall bundle is sold for the same total amount. This strategy allows firms to compete for different segments of the heterogeneous population of consumers

with their differentiated products. Such strategies may be relevant for our market, because the product under consideration may be offered to induce consumers to buy other products along with it or as a substitute for it, so multiproduct strategies may be important in our market and may contribute to heterogeneity in strategies.

2.2. Empirical Literature

Our paper investigates the same market and uses the same data as Ellison and Ellison (2009a, 2009b). The focus of our paper is quite different, though. Ellison and Ellison (2009a) estimates demand in the market, but does not consider dynamic pricing patterns or firm interactions, which is the focus of the current study. Instead, we will use the demand estimates from that paper as an input into our analysis. Ellison and Ellison (2009b) also estimates aspects of demand to see how sensitive consumers are to sales-tax savings from purchasing online and whether consumers exhibit home-state “biases” in their purchasing.

The two empirical papers most closely related to ours both in substance and empirical setting are Edelman and Ostrovsky (2010) and Zhang (2010). They both document the occurrence of Edgeworth cycles in sponsored search, or online advertising, auctions. At first thought, the empirical settings seem quite different because those papers are studying bids in auctions, not prices set by firms in a market. On second thought, though, one can actually draw a very close analogy. In one case, firms are bidding on keywords, where a bid is the amount the advertiser pays for every click-through. Unlike a standard auction, the second, third, etc., highest bidders do not lose out in those auctions, though. They simply are allocated less desirable (*i.e.*, less likely to be clicked) spaces on the results pages. (Typically, the highest bidder is listed first on the page, the second highest second, and so forth.) Also unlike standard auctions, keyword auctions occur continuously in real time and advertisers can change their bids at will, just like a seller can change price in an online market. The analogy to our setting is the following: The price the firm sets to sell its product is like its bid for a favorable rank on the price search engine page and that rank determines the number

of customers it receives, just like an advertiser’s position on the results page determines click-throughs. We do not, in fact, find Edgeworth cycles in our price series, unlike Edelman and Ostrovsky (2010) and Zhang (2010), although our objective of documenting the pricing patterns that do exist is similar.⁵

As mentioned in the introduction, our paper is also related to a literature that tries to characterize interesting movements (and perhaps cycles) in retail gasoline prices, including Castanias and Johnson (1993), Eckert and West (2004), Noel (2007a, 2007b, 2008), Hosken, McMillan, and Taylor (2008), Atkinson (2009), Wang (2009), Lewis (2009), and Doyle, Muehlegger, Samphantharak (2010). Our hourly data is collected at a higher frequency than most of these studies, which mostly use daily data, an exception being Atkinson (2009), which uses bi-hourly observations. Also related is the literature on wholesale gas pricing. Borenstein, Cameron, Gilbert (1997) identify an asymmetry in price responses to cost rises versus falls.⁶ Davis and Hamilton (2004) provide structural estimates of a menu-cost model. They find an adequate fit of the data, but some specific parameter values have implausible sizes, which they use as evidence in favor of alternative models of price inertia. Their preferred alternative is that firms keep prices constant in order to prevent undesirable strategic reactions by customers or rivals. They have a very specific menu-cost model (that of Dixit 1991), which does not nest the inattention story which we will consider as another alternative here. Lewis (2009) shows that temporary shocks may have long-lived strategic effects in the wholesale gas market.

Our finding of price dispersion is related to a much larger literature on this issue. Lach (2002) characterizes price dispersion for a cross-section of retail items sold in Israeli stores. Barron, Taylor, and Umbeck (2004), Hosken, McMillan, and Taylor (2008), and Lewis (2008)

⁵Many advertisers in sponsored search auctions use automated bidding programs, which, combined with the first price auction format, leads to inherently unstable equilibria, according to Edelman and Ostrovsky. Very similar incentives would exist in our market, but our sellers do not use automated price-setting programs, which could be an important reason that we do not find Edgeworth cycles.

⁶A substantial literature, with many contributions by macroeconomists seeking to understand pricing over the business cycle, studies this issue of asymmetric price adjustments to market conditions. See Klenow and Malin (forthcoming) for a survey. In a study of strategies rather than price outcomes, Bhaskar, Machin and Reid (2002) survey over 70 owner-managed Scottish firms to determine how they respond to rivals.

document price dispersion in retail gasoline. Perhaps closest to the present paper among these is Baye, Morgan, and Scholten (2004) because they document price dispersion in an online price comparison site, so a setting similar to ours.

More specifically we are interested not just in heterogeneity in price outcomes—price dispersion—but heterogeneity in strategies used in a dynamic game. Hosken, McMillan, and Taylor (2008) identify strategy heterogeneity in retail gasoline. Lewis (2008) also documents strategy heterogeneity, finding that his stylized facts are hard to reconcile with a single existing theory.

There is a literature, largely methodological at this point, on the structural estimation of dynamic games. The most prominent papers are Aguirregabiria and Mira (2007) and Bajari, Benkard, and Levin (2007). We see our paper as a complement to that literature in the following sense. Those models impose restrictions on firm behavior, that all firms' actions are consistent with Markov perfect equilibria, for instance, to back out certain structural parameters. We want to, instead, remain entirely agnostic about what type of equilibrium strategies (if any) firms are playing and instead focus on the empirical determinants of their actions, resulting markups, and so forth. Imposing those restrictions on firm behavior is neither necessary nor helpful for our exercise.

An experimental literature has investigated markets in which participants can adjust price continuously or nearly so (Millner, Pratt, and Reilly 1990, Davis and Korenok 2009). Another related experimental paper is Leufkens and Peeters (forthcoming), which studies a laboratory market in which prices are set dynamically with various commitment periods across treatments.

3. Empirical Setting

We look at the online market for computer components mediated by a price search engine called Pricewatch. This is an empirical setting examined also in two previous papers, Ellison and Ellison (2009a, 2009b). We focus on aspects of the market and firm behavior that those

papers largely abstracted away from, although details of the empirical setting discussed there will be important here as well. One should consult those papers for a more complete description.

During the period in which our data were collected, the Pricewatch universe was characterized by a large number of small, undifferentiated e-retailers selling memory upgrades, CPUs, and other computer parts. These retailers tended to do little or no advertising, have rudimentary websites, receive no venture capital, and run efficient, profit-maximizing operations. They also tended to receive a large fraction of their customers through Pricewatch. Potential customers could use Pricewatch to locate a product in one of two ways. They could either type a technical product description, such as “Kingston PC2100 512MB,” into a search box, or they could run through a multilayered menu to select one of a number of predefined product categories, *e.g.*, clicking on “System Memory” and then on “PC100 128MB SDRAM DIMM.” In that case, they would receive back a list of products sorted from cheapest to most expensive in a twelve-listings-per-page format. These pre-defined categories may contain as many as 350 listings from 100 different websites. Figure 1 contains the first page of a typical list, that for PC100 128MB memory modules from October 12, 2000.

We will be focusing on the dynamic aspects of firm price-setting strategy, and it is worth noting that the Pricewatch lists exhibit substantial turnover from day to day and even from hour to hour. On average, five of the twenty-four retailers on the first two pages of the above-mentioned list will change their prices on a given day. Each price change typically moves several other retailers up or down one place on the list. Some websites are clearly big players that regularly occupy a position near the top of the Pricewatch list. From time to time one may observe a firm sitting in the first position for a week or more, but there is no rigid hierarchy.

Some of this turnover can be attributed to the technology used by Pricewatch and the various websites at the time. Pricewatch is a database-based system which relies on e-retailers’ updating their own prices in its database. Our impression is that all (or almost all)

of the retailers were setting prices in Pricewatch manually in the time period we study. A typical retailer has dozens or hundreds of products listed in the Pricewatch database, making it impractical to constantly monitor one's place on each predefined product category and the current wholesale prices for each product. Instead, a retailer might manually examine its position on the most important Pricewatch categories a few times a day and might look at current wholesale prices at most once or twice a day.

Alternatively, a sudden drop or uptick in sales (which a retailer might notice if he is monitoring incoming sales in real time) could indicate that the retailer has been bumped up or down in rank. Figure 2, derived from demand estimates from Ellison and Ellison (2009a), indicates how a retailer's daily sales vary with rank. The bulk of sales go to the two or three lowest-priced retailers in this market, but positive sales still accrue to many additional retailers on the list.

4. Data

4.1 Sources

Our primary data source is prices downloaded from Pricewatch.com. We collected information on the twenty-four lowest price offerings, or the first two pages, in the category of 128MB PC100 memory modules.⁷ In addition to the price, we downloaded the specific product name and the firm name, so that we could follow firm pricing strategies for specific products over time.⁸ Due to the frequent turnover on the price lists, it was important to collect these data at high frequency, and, in fact, they were collected hourly from May 2000 to May 2001.⁹

In some calculations, we supplement the pricing data with proprietary data from a retailer who operates three websites listed on Pricewatch. This retailer provided us with daily

⁷We collected price data on several different product categories from Pricewatch. We focus on this particular product category because it is the most active and highest volume category for which we have data.

⁸The products in this category are, physically at least, fairly homogeneous. We do treat it as a change in product offering when a firm changes the name of its product, though.

⁹We used Go!Zilla to carry out the downloads.

wholesale acquisition cost data (which we have reason to believe is common to all retailers). Note that wholesale purchases are typically made every afternoon based on that day’s retail orders, so little or no inventory is held. The retailer also provided us with sales data. We can use those data to calculate margins and markups as well as a proxy for product-level profit by firm conditional on current costs and a firm’s Pricewatch position.

A large number of firms made brief appearances on the Pricewatch lists. Since we are interested in the dynamics of firms’ pricing patterns, we only retained firms that were present for at least 1,000 hours at some point during the year (approximately one-eighth of our time period) and changed price while staying on the list at least once. We also deleted the small number of firms who had multiple products on the first two pages of Pricewatch simultaneously. We were left with 43 firms that appear at some point during the year, with at most 24 present at any particular moment.

4.2 Variables

Based on these data, we created a number of variables to describe factors that might be important to firms’ decisions about timing and magnitude of price changes. The rank of that firm on the Pricewatch list is an obvious candidate, but so would be factors like markup, length of time since its last price change, number of times a firm has been “bumped,” or had its rank changed involuntarily, since its last price change, and so forth. Table 1 contains a description of these variables, and Table 2 lists their summary statistics.

Definitions of a few of the variables could use additional explanation. *Placement* is simply the fraction of the distance between the next lower-priced firm and the next higher-priced firm a particular firm is in price space. In other words, if three consecutive firms were charging \$85, \$86, and \$88, the value for *Placement* for the middle firm would be 0.33, or $(86 - 85)/(88 - 85)$. *Density* is a measure of how crowded together firms are in price space in the immediate region around a particular firm. It is defined as the difference between the price of the next higher-priced firm minus the price of the firm three spaces below divided

by 4. In other words, if five consecutive firms were charging \$84, \$84, \$85, \$86, and \$88, the value for *Density* for the firm charging \$86 would be 1, or $(88 - 84)/4$. *CostTrend* and *CostVol* were computed by regressing the previous two weeks of costs on a time trend and using the estimated coefficient as a measure of the trend and the square root of the estimated error variance as a measure of the volatility. The definitions of the remaining variables are self-explanatory.

Turning to Table 2, note first that the average price charged by our firms for a PC100 128MB memory module during this period was about \$69 with a very large range, \$21 up to \$131. Most of this variation occurs over time, with prices typically above \$100 at the beginning of the period and down in the \$20s by the end, mirroring a large decline in the wholesale cost of these modules.

4.3 Markups

We will use markups throughout the paper as a handy indicator of firm strategies and market outcomes. It is useful at this point to look more closely at this measure. As noted in Table 1, we compute markups as the difference of listed price and wholesale acquisition cost of the memory module divided by the listed price, or $\frac{P-C}{P}$. Like the accounting “gross margin,” this measure does not reflect any fixed costs or other costs of order fulfillment and therefore overstates the per-item profit that the retailer enjoys. Despite that, the summary statistics indicate that the markups here are extremely low, on average 4% but some as low as –35%! The explanation for the very low, often negative, markups is an add-on pricing policy employed by the retailers, which is the main focus of Ellison and Ellison (2009a). Since, to the best of our knowledge, all of the firms in the market were engaging in these policies during this period, it is still useful to examine relative markups across retailers and over time, even if the markups themselves suggest lower profits than these firms were actually making.

In addition to the markups being low, they display a surprising amount of variation,

ranging from -35% to 38% . Figure 3 is a histogram of the hourly markups of retailers throughout the year. The bulk of the markups are small and positive, and the histogram exhibits skew, with the positive tail being much thicker than the negative tail. Figure 4, a plot of markups over time for particular ranks, helps decompose the sources of this substantial variation. One thing is clear: markups are not constant over time. In particular, retailers are not simply choosing a desired markup and changing price daily to maintain it as costs change. Instead, this figure suggests that markups are highly volatile within retailer over time. The primary explanation is the underlying volatility of costs combined with periods of inattention on the part of the price-setters. For instance, retailers might check the price of their 128mb PC100 memory modules every few days or once a week, during which time volatile costs could have changed their markups substantially. Note also that the quantity-weighted markups are fairly close to the rank 1 markups. This proximity is simply a result of sales being heavily weighted toward the lowest-priced firms. We also computed the autocorrelations of markups by product. The hour-to-hour autocorrelations are high, ranging from 0.95 for some products to greater than 0.99 for others. Since costs only change once a day and most retailers do not change prices most days, these figures are not surprising. The day-to-day autocorrelations, however, exhibit much more range, with some products at 0.85 and others as low as 0.20. This range suggests a heterogeneity in pricing strategies, where some firms maintain fairly stable markups and others do not.

5. Classifying Firms' Strategies

One of the striking characteristics of our data set was heterogeneity in firm strategies. Any empirical model of pricing behavior would have to accommodate this heterogeneity. Modeling pricing on a firm-by-firm basis would, of course, allow for arbitrary firm-level heterogeneity, but that analysis would sacrifice power, perhaps unnecessarily, and, more importantly, would result in a non-random sample selection, eliminating almost all of the less active firms due to too few observations. To balance those two concerns, we decided to categorize firms into

classes, or what we will call “strategic clusters,” as a prelude to empirically modeling their pricing behavior. To do so, we use the technique of cluster analysis.

Cluster analysis is a well-known tool in many other social sciences and sciences but is relatively unknown in empirical economics.¹⁰ Therefore, a concise description makes sense here. We should begin by saying that we think of cluster analysis more as a handy tool to organize data as opposed to a statistical technique which produces useful estimates of anything or is optimal in any sense. With that in mind, we would define cluster analysis as a set of techniques used to assign objects, which are differentiated on multiple dimensions, to clusters which contain other, similar objects. In our case, the objects are firms, and we seek to assign them to clusters based on their strategic actions, such as how often they change their price, where on the price-sorted lists they want to be, whether they change as a result of being bumped from their rank, and so forth.¹¹ We think allowing for three clusters of firms is a nice balance between spanning most of the important firm heterogeneity we want to capture and still obtaining clusters large enough to analyze empirically with some confidence.

We started by standardizing the variables describing firm strategy to give them all a standard deviation of one. This step is common prior to performing a cluster analysis to prevent differences in the variable with the largest variance from dominating the assignment. Then we defined similarity (or dissimilarity) between two firms as the Euclidean distance in the space defined by the variables. Finally, we chose a method which defines similarity between clusters as a function of the sum of squared differences between all firms in the two clusters, Ward’s method.

We performed an agglomerative cluster analysis, which first places every firm in its own

¹⁰A standard reference for cluster analysis is Romesburg (2004).

¹¹We used seven variables for the cluster analysis. The averages at the firm level of targeted *Rank* and targeted *Placement* were used, in other words the *Rank* and *Placement* that firms had immediately after a price change, averaged within each firm. In addition, we used firm-level average of *NumBump* and of *SinceChange* and firm-level variance of *SinceChange*. Finally, we used a firm’s total time present on the list and the amount of time the firm spent in ranks 1-12 as a fraction of its time present on the list. Our goal was to cluster on a group of variables over which the firm should have a fair amount of control as opposed to variables that describe outcomes that they experience.

cluster and proceeds iteratively by agglomerating the most similar clusters until, in our case, there are three. We used the standard Stata cluster analysis command with the options described above. Table 3 contains the output of the clustering, both in terms of the number of firms which are assigned to each of three clusters and the average values of key variables within clusters. These averages provide an intuitive description, or shorthand, for each of the clusters. Cluster 1, for instance, could be characterized as the cluster of active firms who like to populate the highest spots (lowest ranks) on the price-sorted lists. These firms change prices often, tend to have low markups, and tend to not stray far from their preferred locations on the price-sorted lists. Cluster 2, the smallest cluster, is made up of firms with moderate prices and markups who are relatively inactive. In other words, these firms may target middle ranks but tend to get bumped far from their preferred spots before they bother to change their prices. Cluster 3 is higher-priced but more active than cluster 2. They prefer lower spots (higher ranks) on the price-sorted list and will change prices relatively often to maintain their preferred locations. To summarize, cluster 1 is low-priced and active, cluster 2 is mid-priced and inactive, and cluster 3 is high-priced and fairly active.

This characterization is nicely illustrated in Figure 5. We chose one representative firm from each cluster and graphed both their prices and ranks for the month of August, 2000. Note that that month, like most of our months, is marked by falling wholesale prices, so a firm listing a particular price would tend to get bumped higher in rank as firms around it were lowering their prices in response to falling costs. Just as the means by cluster suggest, our representative firm from cluster 2 is between the lower price of the cluster 1 firm and the higher price of the cluster 3 firm. In addition, it is noticeably less active than either of its counterparts, changing price only twice during the month. As a result of its inaction, it is bumped up from a rank of 6 at the beginning of the month¹² to 14 by the end. Both the cluster 1 and cluster 3 firms are close to their initial rank at the end of the month. There is a notable difference between the firms from clusters 1 and 3 other than price level, though.

¹²It only enters on August 5 and is not present during the first four days of the month. Likewise, the representative cluster 3 firm enters about halfway through the month, August 14.

The firm in cluster 3 is displaced further from its initial rank over the course of the month than the firm in cluster 1. It does, though, tend to regain its initial rank when it is bumped too far, unlike the firm from cluster 2.

An observant reader might notice phenomena in Figure 5 that resemble Edgeworth cycles, or in fact, upside-down Edgeworth cycles. Recall an Edgeworth cycle involves firms gradually decreasing prices to slightly undercut rivals and then reverting back to a high (monopoly) price after marginal cost has been reached. Although our cycles might bear some resemblance to Edgeworth cycles, their origin is entirely different. First, they are cycles in rank, not price, and rank is not something that firms can directly control. Second, our cycles result from inactivity, not the series of small price changes that give rise to Edgeworth cycles. Firms get bumped away from their desired rank—if cost is falling, they would typically be bumped to higher ranks—until they are sufficiently far to warrant the effort to change price and regain a desirable rank.

6. Modeling Firms' Price Changes

We use a two-stage model of firms' price-setting behavior. Each period t , corresponding to an hour in the dataset, the first stage involves each firm deciding whether to keep the same price from period $t - 1$ or to change it. Over stretches of time, either because market conditions or rival actions have remained stable, or because the firm does not wish to exert effort of attending to current conditions just yet, the firm will decide not to change price. If it decides to change price in a certain period t , however, it moves to a second-stage decision about the size of the price change, including the direction of change (increase or decrease). We estimate the first-stage decision—timing of price change—using a probit specification, in which the dependent variable is an indicator equalling 1 if firm i changes its price between period $t - 1$ and t and 0 if firm i keeps price the same. The results are reported in Table 4. We estimate the second-stage decision—size of price change—using an ordinary least squares regression in which the dependent variable is the proportional change in price: $\ln Price_{i,t} - \ln Price_{i,t-1}$.

The results are reported in Table 5.

We estimate three separate models for each of the three strategic clusters of firms, thus allowing us to characterize the heterogeneous strategies across clusters and to reflect this heterogeneity in the simulations. In addition to these cluster-specific models, we present an initial model that combines all the firms in one model, as a sort of general summary of the “average” firm strategy. We present two specifications of each model. One is a parsimonious specification that includes a small set of the most important variables which we believed to have a first-order effect on firms’ pricing decisions: variables reflecting price, such as *Markup*, variables capturing characteristics of the current spell since the last move, *SinceChange* and *NumBump*, and variables describing recent cost movements, *CostTrend* and *CostVol*. This is the specification that will be used to generate the simulations below, the parsimony helping to generate stable simulations which turn out to be difficult with the inclusion of less-important variables. To help provide a full characterization of firms’ strategies from a descriptive perspective, we also estimate a rich specification that includes additional variables which may potentially affect firms’ pricing, albeit in a less powerful way. The rich specification includes controls for the distribution of prices in the neighborhood of firm i (*Placement* and *Density*) and aspects of distribution of prices more globally (*LowPrice*, *PriceRange*, and *AvgMarkup*).

6.1. Timing of Price Change

Consider the results for the probit model of the timing of price change reported in Table 4. All results are presented as marginal effects for ease of interpretation, and all but the results for the markup variables (*Markup* and *Markup*²) are scaled up by 10^3 for legibility. The combined probit involves over 165,000 observations, over 60% of which are for firms in strategic cluster 1 (the preponderance of observations due to the fact that half the firms are in cluster 1 and these firms tend to have more complete presence during the sample period). The combined probits in the first set of columns will generally be somewhere in the middle

of the three clusters' results, but somewhat skewed toward the cluster 1 results. Reported standard errors are robust to heteroskedasticity and are adjusted to account for the possible non-independence across observations for the same firm.¹³

It should be noted that relatively few of the results for cluster 2 are statistically significant. This is true for two reasons. First, one of the central characteristics of the cluster is to adjust prices relatively infrequently. Hence there will not be much variation in the dependent variable to explain. Second, although the probits for cluster 2 include almost 20,000 observations, this is still only a small fraction of the overall sample. This is in fact a small sample once one recognizes that time-series observations for a single firm are not independent and that there are few price changes (showing up as 1's for the dependent variable) in the subsample. Overall, the pseudo- R^2 values are fairly low across the probits, only as high as 0.034 (i.e., 3.4%), evincing the difficulty in predicting the rare event of a price change in hourly data.

The marginal effect of *Markup* in the probits is difficult to discern directly because it enters as a quadratic to allow for possible nonlinear effects. To aid in interpretation, Figure 6 graphs the quadratic function associated with the marginal effects from the parsimonious specification. The dotted curve for the combined results, in particular its steep downward slope for negative margins, indicates that the average firm is very sensitive to large negative margins and likely to adjust their prices then. For example, a price-cost margin of -0.35 , meaning that price is 35% below cost, leads the average firm to increase its probability of a price change in a given period by almost 0.012, that is, 1.2 percentage points. This may seem like a small increase in probability until one recalls that this is hourly data, and the probability of a price change in any given hour is quite low to begin with. The hourly marginal effect cumulates into a larger effect over longer time periods, implying that a 1.2 percentage point marginal effect is actually quite large. The effects for cluster 1 firms are

¹³These are sometimes referred to as clustered standard errors, a usage of "cluster" which should be kept distinct from the strategic clusters used throughout this paper. Clustered standard errors adjust for the correlation among errors for a single firm, whereas a strategic cluster involves a group of firms identified by the methods of cluster analysis.

even stronger, with a markup of -0.35 leading to a 0.013 increase in the probability of price change in an hour. Clusters 2 and 3 show some but much less sensitivity, reflecting the more general result that cluster 1 includes more responsive firms than the other clusters.

The fact that the curve for cluster 3 in Figure 6 passes through the origin and becomes negative for positive margins indicates that positive margins lead firms in this cluster to be slower to adjust prices. Evidently, they are willing to trade off increased margins for lower sales in this region. The curves for clusters 1 and 2 tell a different story. Positive margins less than 30% have little effect on the firms' likelihood of changing prices, as can be seen from the curves hovering close to the horizontal axis. For large, positive margins, cluster 1's curve rises steeply and becomes positive, and cluster 2's curve also becomes positive, indicating that the largest positive margins also lead these firms to adjust prices, presumably because they are losing demand. For example, projecting out to a margin of 0.4 (price 40% above cost, which is slightly outside our data range), this margin would lead a cluster 1 firm to be 0.003 (0.3 percentage points) more likely to change price in a given hour.

The *SinceChange* variable is included to determine whether prices are adjusted according to a systematic schedule. Changing prices according to a systematic schedule, say every x periods, would show up as a positive coefficient on *SinceChange* as the firm would be increasingly likely to change price as the next scheduled date for changing price approached. Instead, we find a negative and statistically significant result for the combined and cluster 1 results and results that are still negative but smaller in magnitude for cluster 3. For these clusters, a firm's price changes tend to be bunched in time, followed by longer spells of inactivity than would be predicted by regular schedule. This finding supports an inattention model, but of a certain type, with the manager of each of these firms being drawn away from watching the market at random times for random durations, but then making rapid adjustments in periods he or she is allowed to attend to the market, as opposed to setting prices and then returning to revise them at consistent intervals, say once every other day. Cluster 2's results are essentially 0, indicating that pricing for these firms may conform more

to a regular schedule. The examples in the lower panels of Figure 5 illustrate this pattern, with the representative firm from cluster 2 showing price changes at regular intervals but those for the other clusters being more scatter-shot.

Since *NumBump* is the net number of ranks a firm has been bumped, positive or negative, we included the absolute value as an explanatory variable, reasoning that being bumped a long distance in either direction from your desired rank could precipitate a price change. $|NumBump|$ is, in fact, significant in all specifications and positive as expected: the further you are bumped in either direction from the rank you last chose, the more likely you are to change your price.

Finally, movements of wholesale costs appear to be an important determinant of when to change prices, but in different ways for the three clusters. Cluster 1 is moved to change price either when costs are rising quickly or falling quickly, cluster 3 when costs are falling quickly, and cluster 2 in neither case. Cluster 3 firms, who tend to have high ranks, are sensitive to costs falling because, if they do not act, they will soon find themselves out of the lowest 24 prices as firms come in and undercut them. Cluster 2 firms are somewhat insulated both from large negative margins and from being bumped from the lowest 24 and can afford to be less sensitive to the vicissitudes of the wholesale market. Cluster 1 firms, in addition to being most active, seem to also be the most sensitive to market conditions. Cost volatility is less important in determining time of price change, although there is some evidence that cluster 3 firms are more likely to change price during periods of high cost volatility.

It is interesting to note the heterogeneity across clusters in the importance of the remaining variables as well. For instance, *Placement* has the expected negative sign in cluster 2, indicating that a firm that is well positioned relative to its neighbors in price space is less likely to want to change its price. The ideal placement (associated with the highest value of *Placement*) is one just below the higher-priced neighbor because this maximizes markup subject to keeping the firm's rank constant, and rank is an important determinant of demand in this market. This effect is not significantly present for the other clusters, though. *Density*

is significant and has the expected positive sign for clusters 1 and 2, where firms are more likely to adjust prices if their price is close to a concentration of other firms', but is not important for cluster 3.

6.2. Size of Price Change

Turn next to the results for the size of price change reported in Table 5. These ordinary least squares regressions were run on only the subset of observations for which there was a price change from the previous hour. This results in a much smaller sample because price changed in fewer than 1% of the observations. The cluster 2 regressions should be interpreted with particular caution because they only include 70 observations.

The dependent variable is the change in log price, $\ln(\text{Price}_{it}) - \ln(\text{Price}_{it-1})$, so effectively a percentage change for small price changes. As before, the reported standard errors are heteroskedasticity robust, adjusted to account for non-independence across observations for the same firm. Besides the first two variables (constant and *Markup*), the coefficients for the remaining variables are scaled by 10^2 for legibility. The results are generally sensible and consistent with prior expectations. Notice that the R^2 in Table 5 is an order of magnitude higher than the pseudo R^2 from Table 4, implying that it is much easier to predict the size of price change than its exact timing at an hourly frequency.

We will focus on the more robust, substantial, and significant results, starting with the constant term. Comparing the constant term across clusters 1 and 3, we see that cluster 1 firms on average have larger reductions in price when they change it compared to cluster 3 firms, but cluster 2 has the largest. This is consistent with cluster 1 maintaining low prices and cluster 2 having to make large moves because they are infrequent. *Markup* is negative and significant across all columns. This implies that firms adjust prices to re-establish their ideal margins, increasing price when margins are negative and decreasing them when margins become too large.

We created two new variables based on *SinceChange*, which are *SinceChange* interacted

with dummies for whether cost had been rising or falling. We consider these variables because, although we would expect firms to make large price changes after long periods of inactivity, those changes could be large and positive or large and negative. In other words, we might imagine that *SinceChange* on its own would be more likely to increase error variance than to predict whether a price change would be positive or negative. The interactions allow *SinceChange* to have differential effects depending on whether costs, and therefore prices, are trending up or down. Despite our reasoning, though, the *SinceChange* variables seem to be wholly unimportant. *NumBump*, now entering as-is instead of as its absolute value, is negative and highly significant in all specifications—the net number of spaces you have been bumped is an important determinant of how far you move conditional on moving. Magnitudes are also fairly similar across clusters.

Unlike in the previous regressions, we do not have *a priori* reasons to separate out potential effects of positive and negative cost trends since we expect monotonicity—when *CostTrend* is positive, we expect most price changes will be positive and when it is negative, we expect most price changes will be negative. Also unlike the previous regressions, *CostTrend* does not appear important here. *CostVol*, with the exception of the combined results, is also not significant.

It is interesting to note one broad difference between the results from the two different empirical models, Table 4, timing of change, and Table 5, size of change. Heterogeneity across clusters stood out as an important feature in Table 4. It was hard to miss that different factors were driving the different clusters to act. In Table 5, however, this heterogeneity was much less apparent. Conditional on deciding to act, similar factors drive the firms' actions.

7. Simulations

In the early 1980s, the political scientist Robert Axelrod carried out a series of simulations which changed the way that social scientists thought about repeated games (Axelrod, 1984). Axelrod invited game theory colleagues to submit strategies for a simple two-person repeated

game. He then pitted the strategies against one another in a tournament of simulations to see which strategies did best at maintaining prices closest to the monopoly price against themselves, different strategies, the entire field of strategies, and so forth. A simple strategy called “tit-for-tat,” where firms punish their rivals’ one-period deviations from the monopoly price by their own one-period deviations, performed the best (resulted in the highest markups) against a wide range of other strategies. The simulations we perform here are in this spirit: we can modify strategies that firms in this market actually use and pit those modified strategies against one another to simulate the market outcomes under these modifications. For instance, how would market outcomes change if, say, everyone was using cluster 1 strategies?

Since Axelrod’s game was entirely an invention, he was able to control all aspects of it. In particular, he specified cost functions and profit functions for firms, assuring homogeneity in those dimensions. Since our simulations are based on actual firms operating in an actual market, we have no such control, but we can also structure our simulations based on quantities we can observe and estimate in our market, which Axelrod could not do. We create a set of twenty-four simulated firms and assign them starting markups by drawing randomly from the empirical markup distribution we observe. From those markups, we then can back out prices (given cost) and generate ranks. We then allow cost to move as it actually did and the simulated firms’ prices to evolve as governed by the empirical models we developed and discussed in the previous section. We can alter the mix of firms competing at will, by simply having different proportions of them governed by the three cluster-level empirical models. For instance, to simulate the dynamics of price competition with a group of firms entirely composed of active, low-price firms, we just have all of our simulated firms’ hour-by-hour decisions arising from cluster 1 empirical estimates.

Note that to ensure well-behaved simulations, we use the more parsimonious specifications for both the probability-of-change and size-of-change models. In particular, we allow the variables *Markup*, *SinceChange*, and *NumBump*, as well as firms sensitivity of costs movements, to drive the firms’ decisions on timing and size of price change in the simulations.

Also, we allow for entry of new firms in the simulations. Firms enter at random times and random ranks on the list consistent with the frequency and location of firm entry in the actual market.¹⁴ Note that firm entry has the effect of knocking the highest-priced firm off the list so that we always maintain twenty-four firms.

We performed four different simulations. Our first was simply a simulation where we reproduced the actual mix of the three strategic clusters we observe.¹⁵ The second simulation imposed the counterfactual that all firms are of cluster 1. The third simulation imposed the counterfactual that all firms are of cluster 2. Finally, the fourth simulation imposed the counterfactual that all firms are of cluster 3.

Table 6 contains the results from our simulations.¹⁶ Note first that our simulations maintaining the actual mix produce slightly smaller markups than the actual markups. For all firms, markups averaged 3.43%, but our simulation produced markups of 2%. The ordering of markups by class was maintained, though, with class 1 having the lowest markup and class 3 the highest. The last three columns in the table are the average markups from our counterfactual simulations. It is interesting to note how the average markups change when only one type of firm is present. For cluster 1, the markups edge down slightly, from -0.4% to -1.4% . Strikingly, the cluster 2 markups change quite substantially, going from 4% to almost 20% when only firms of that type are present. Finally, the cluster 3 markups also actually edge down slightly but stay close to 4% . Recall that we characterized cluster 2 firms

¹⁴For class 1, firms entered with probability $8/720$ each hour, which resulted in eight expected entrants per month. They entered uniformly over rank. For class 2, firms entered with the probability $12/720$ each hour and were twice as likely to enter ranks 15-24 as ranks 1-14. For class 3, firms entered with probability $16/720$ each hour and were twice as likely to enter ranks 15-20 and three times as likely to enter ranks 21-24 as ranks 1-14. These figures were chosen to approximate actual entry rates and locations of firms of different classes.

¹⁵The mix of firm types varies over the year, of course, so we used a mix based on the fraction of our 43 firms that are assigned to each of the three strategic clusters. For 24 firms, that mix turns out to be twelve in cluster 1, five in cluster 2, and seven in cluster 3.

¹⁶For each of the four scenarios we explored through simulations, we performed five simulated years. The range in average markups across the five years was quite small in each case, -0.0179 to -0.0114 for the cluster 1 simulations, for instance. We report here the average markups averaged across the five simulated years. Simulating a large number of years in each case would have been impossible given our current computing constraints, but even if it had been possible, we are quite confident that the results would be largely unchanged.

as mid-priced but least active, so it is perhaps not surprising that the counterfactual with all firms of that less-active type would result in substantially higher markups than cluster 2 firms had in the actual-mix simulation. Note, though, that cluster 2 was the smallest cluster, so the counterfactual involving all firms being in that cluster is, in some sense, furthest from the mix we observe. Cluster 3, the moderately active and highest-priced firms have, not surprisingly, the highest average markups in the actual-mix simulation, and are able to maintain those in the cluster-3-only simulation.

In addition to the summary statistics on simulated markups, we might also be interested in the distribution and dynamics of these markups. In Figure 7, we graph the markups over time for the median firm from each of our four simulations and include actual median markups for comparison. It is notable how similar the dynamic patterns are with the striking exception of the cluster 2 simulations. Recall that cluster 2 firms' probability of changing price was estimated to be roughly linear in markup while the other clusters had U-shaped relationships. This lack of a disciplining effect on extreme markups could give rise to this gradual increase in markups we see as cost goes down. It is important to note, though, that cluster 2 firms did not wander off into large markups in our actual mix simulation. If their lack of sensitivity to large markups did not offer enough discipline to keep them in line when only they were present, the presence of cluster 1 and 3 firms did.

Finally, in Figure 8, we produce a simulated counterpart to our Figure 5, the graph showing rank and price movements for three representative firms. The simulated firms appear to be somewhat more active and engage in larger price changes compared to the very modest changes we see in Figure 5. The cycles in rank, especially for the cluster 2 firm, are still evident, though.

We view this set of simulations as only a first attempt at using our methodology to examine how the dynamics of firm interactions can change when various aspects of the setting, such as the mix of firm types, changes. There are a variety of other simulation exercises that we can perform using this framework. For instance, we could hold fixed firm

strategies and vary market conditions, such as degree of cost volatility or the lags with which firms can observe rivals prices. Or we could hold firm strategies mostly fixed but just change one aspect, say, the degree of sensitivity to ones rank.

8. Conclusion

The empirical setting and particular data we have gathered lend themselves to the study of repeated-game strategies. The observed patterns in prices, price ranks, and markups have not been well documented to date, and part of our objective is their careful description. In addition to simply documenting these patterns, we develop empirical models of firms' price-changing behavior, and we construct a framework for simulating market outcomes when aspects of the empirical setting are altered. We view this framework as quite flexible, allowing many types of simulations beyond the preliminary ones we present here.

References

- Aguirregabiria, Victor and Pedro Mira. (2007) "Sequential Estimation of Dynamic Discrete Games," *Econometrica* 75: 1-53.
- Atkinson, Benjamin. (2009) "Retail Gasoline Price Cycles: Evidence from Guelph, Ontario Using Bi-Hourly, Station-Specific Retail Price Data," *Energy Journal* 30: 85–110.
- Axelrod, Robert. (1984) *The Evolution of Cooperation*. New York: Basic Books.
- Bajari, Patrick, C. Lanier Benkard, and Jonathan Levin. (2007) "Estimating Dynamic Models of Imperfect Competition," *Econometrica* 75: 1331-1370.
- Barron, John M., Beck A. Taylor, and John R. Umbeck. (2004) "Number of Sellers, Average Prices, and Price Dispersion," *International Journal of Industrial Organization* 22: 1041–1066.
- Baye, Michael R., John Morgan, and Patrick Scholten. (2004) "Price Dispersion in the Small and in the Large: Evidence from an Internet Price Comparison Site," *Journal of Industrial Economics* 52: 463–496.
- Bhaskar, V., S. Machin, and G. Reid. (1991) "Testing a Model of the Kinked Demand Curve," *Journal of Industrial Economics* 39: 241–254.
- Bhaskar, V. and Fernando Vega-Redondo. (2002) "Asynchronous Choice and Markov Equilibria," *Journal of Economic Theory* 103: 334–350.
- Borenstein, Severin, A. Colin Cameron, and Richard Gilbert. (1997) "Do Gasoline Prices Respond Asymmetrically to Crude Oil Price Changes?" *Quarterly Journal of Economics* 112: 305–339.
- Cahuc, Pierre and Hubert Kempf. (1999) "Asynchronized Multiperiod Commitments and Cycles," *Journal of Economic Behavior and Organization* 40: 387–407.
- Castanias, Rick and Herb Johnson. (1993) "Gas Wars: Retail Gasoline Price Fluctuations," *Review of Economics and Statistics* 75: 171–174.
- Chowdhury, Prabal Roy. (2008) "Bertrand-Edgeworth Equilibrium with a Large Number of Firms," *International Journal of Industrial Organization* 26: 746–761.
- Davis, Douglas D. and Oleg Korenok. (2009) "Posted Offer Markets in Near-Continuous Time: An Experimental Investigation," *Economic Inquiry* 47: 449–466.
- Davis, Michael C. and James D. Hamilton. (2004) "Why Are Prices Sticky? The Dynamics of Wholesale Gasoline Prices," *Journal of Money, Credit and Banking* 36: 17–37.
- Dixit, Avinish. (1991) "Analytical Approximations in Models of Hysteresis," *Review of Economic Studies* 58: 141–151.

- Doyle, Joseph, Erich Muehlegger, and Krislert Samphantharak. (2010) “Edgeworth Cycles Revisited,” *Energy Economics* 32: 651–660.
- Eckert, Andrew. (2003) “Retail Price Cycles and the Presence of Small Firms,” *International Journal of Industrial Organization* 21: 151–170.
- Eckert, Andrew. (2004) “An Alternating-Move Price-Setting Duopoly Model with Stochastic Costs,” *International Journal of Industrial Organization* 22: 997–1015.
- Eckert, Andrew and Douglass West. (2004) “Retail Gasoline Price Cycles across Spatially Dispersed Gasoline Stations,” *Journal of Law and Economics* 47: 245–271.
- Edelman, Benjamin and Michael Ostrovsky. (2010) “Strategic Bidder Behavior in Sponsored Search Auctions,” *Decision Support Systems*, forthcoming.
- Edgeworth, Francis. (1925) “The Pure Theory of Monopoly,” in *Papers Relating to Political Economy*, vol. 1. London: MacMillan, 111–142.
- Ellison, Glenn. (1994) “Theories of Cartel Stability and the Joint Executive Committee,” *RAND Journal of Economics* 25: 37–57.
- Ellison, Glenn and Sara Fisher Ellison. (2009a) “Search, Obfuscation, and Price Elasticities on the Internet,” *Econometrica* 77: 427–452.
- Ellison, Glenn and Sara Fisher Ellison. (2009b) “Tax Sensitivity and Home State Preferences in Internet Purchasing,” *American Economic Journal: Economic Policy* 1: 53–71.
- Fishman, Arthur. (1992) “Search Technology, Staggered Price-Setting, and Price Dispersion,” *American Economic Review* 82: 287–298.
- Fudenberg, Drew and Eric Maskin. (1986) “The Folk Theorem in Repeated Games with Discounting or with Incomplete Information,” *Econometrica* 54: 533–554.
- Genesove, David and Wallace P. Mullin. (2001) “Rules, Communication, and Collusion: Narrative Evidence from the Sugar Institute Case,” *American Economic Review* 91: 379–398.
- Green, Edward J. and Robert H. Porter. (1984) “Noncooperative Collusion under Imperfect Price Information,” *Econometrica* 52: 87–100.
- Hansen, Per Svejstrup, *et al.* (1996) “Asymmetric Adjustment in Symmetric Duopoly,” *Economic Letters* 53: 183–188.
- Hosken, Daniel S., Robert S. McMillan, and Christopher T. Taylor. (2008) “Retail Gasoline Pricing: What Do We Know?” *International Journal of Industrial Organization* 26: 1425–1436.
- Klenow, Peter J. and Benjamin A. Malin. (forthcoming) “Microeconomic Evidence on Price-Setting,” in Benjamin Friedman and Michael Woodford, eds., *Handbook of Monetary Economics*, vol. 3. Amsterdam: North-Holland.

- Lach, Saul. (2002) “Existence and Persistence of Price Dispersion: An Empirical Analysis,” *Review of Economics and Statistics* 84: 433–444.
- Lal, Rajiv and Carmen Matutes. (1989) “Price Competition in Multimarket Duopolies,” *Rand Journal of Economics* 20: 516–537.
- Lau, Sau-Him Paul. (2001) “Aggregate Pattern of Time-dependent Adjustment Rules, II: Strategic Complementarity and Endogenous Nonsynchronization,” *Journal of Economic Theory* 98: 199–231.
- Leufkens, Kasper and Ronald Peeters. (2008) “Intertemporal Price Competition with Exogenous Demand Shocks,” *Economic Letters* 99: 301–303.
- Leufkens, Kasper and Ronald Peeters. (forthcoming) “Price Dynamics and Collusion under Short-Run Price Commitments,” *International Journal of Industrial Organization*.
- Lewis, Matthew. (2008) “Price Dispersion and Competition with Differentiated Sellers,” *Journal of Industrial Economics* 56: 654–678.
- Lewis, Matthew. (2009) “Temporary Wholesale Gasoline Price Spikes Have Long-Lasting Retail Effects: The Aftermath of Hurricane Rita,” *Journal of Law and Economics* 52: 581–605.
- Maskin, Eric and Jean Tirole. (1988) “A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles,” *Econometrica* 56: 571–599.
- Millner, Edward L., Michael D. Pratt, and Robert J. Reilly. (1990) “Contestability in Real-Time Experimental Flow Markets,” *Rand Journal of Economics* 21: 584–599.
- Noel, Michael. (2007a) “Edgeworth Price Cycles, Cost-Based Pricing, and Sticky Pricing in Retail Gasoline Markets,” *Review of Economics and Statistics* 89: 324–334.
- Noel, Michael. (2007b) “Edgeworth Price Cycles: Evidence from the Toronto Retail Gasoline Market,” *Journal of Industrial Economics* 55: 69–92.
- Noel, Michael. (2008) “Edgeworth Price Cycles and Focal Prices: Computational Dynamic Markov Equilibria,” *Journal of Economics and Management Strategy* 17: 345–377.
- Porter, Robert. (1983) “A Study of Cartel Stability: The Joint Executive Committee, 1880–1886,” *Bell Journal of Economics* 15: 301–314.
- Romesburg, H. Charles. (2004) *Cluster Analysis for Researchers*. North Carolina: Lulu Press.
- Rotemberg, Julio J. and Garth Saloner. (1986) “A Supergame-Theoretic Model of Price Wars During Booms,” *American Economic Review* 76: 390–407.
- Scott Morton, Fiona. (1997) “Entry and Predation: British Shipping Cartels 1879–1929,” *Journal of Economics and Management Strategy* 6: 679–724.

- Tirole, Jean. (1988) *The Theory of Industrial Organization*. Cambridge, Massachusetts: MIT Press.
- Wallner, Klaus. (1999) “Sequential Moves and Tacit Collusion: Reaction-Function Cycles in a Finite Pricing Duopoly,” *Journal of Economic Theory* 84: 251–267.
- Wang, Zhongmin. (2009) “(Mixed) Strategy in Oligopoly Pricing: Evidence from Gasoline Price Cycles Before and Under a Timing Regulation,” *Journal of Political Economy* 117: 987–1030.
- Zhang, Xiaoquan. (2010) “Finding Edgeworth Cycles in Online Advertising Auctions,” Sloan School of Management working paper.

Table 1: Definition of Variables

Variable	Indexes varies over	Definition
<i>Price</i>	<i>i, t</i>	Current listed price in dollars
<i>Markup</i>	<i>i, t</i>	Proportional markup of price over wholesale cost, $(P - C)/P$
<i>Rank</i>	<i>i, t</i>	Rank of listing in price-sorted order
<i>SinceChange</i>	<i>i, t</i>	Hours since firm last changed its price
<i>NumBump</i>	<i>i, t</i>	Net number of ranks bumped since last price change
<i>Placement</i>	<i>i, t</i>	Placement between adjacent firms in price space
<i>Density</i>	<i>i, t</i>	Measure of density in price space of firms with nearby ranks
<i>FirstPage</i>	<i>i, t</i>	Indicates appearance on first page (rank 1–12)
<i>PriceLow</i>	<i>t</i>	Lowest of currently posted prices
<i>PriceRange</i>	<i>t</i>	Range across currently posted prices
<i>AvgMarkup</i>	<i>t</i>	Average across firms' price margins in levels
<i>NewFirm</i>	<i>t</i>	Indicates new entry into ranks 1–24 in past two weeks
<i>CostTrend</i>	<i>t</i>	Wholesale cost trend over previous two weeks
<i>CostVol</i>	<i>t</i>	Wholesale cost volatility over previous two weeks

Notes: *i* indexes firms and *t* indexes hours.

Table 2: Descriptive Statistics

Variable	Mean	Std. dev.	Min.	Max.	Obs.
<i>Price</i>	68.94	34.71	21	131	165,955
<i>Markup</i>	0.03	0.10	-0.35	0.38	165,955
<i>Rank</i>	11.65	6.73	1	24	165,955
<i>SinceChange</i>	111.97	135.41	1	1,113	165,955
<i>NumBump</i>	1.38	3.62	-23	21	165,955
<i>Placement</i>	0.60	0.42	0	1	165,955
<i>Density</i>	0.56	0.39	0	3	165,955
<i>FirstPage</i>	0.54	0.50	0	1	165,955
<i>PriceLow</i>	62.63	32.90	21	122	8,218
<i>PriceRange</i>	12.62	4.36	5	29	8,218
<i>AvgMarkup</i>	0.04	0.07	-0.12	0.28	8,218
<i>NewFirm</i>	0.63	0.48	0	1	8,218
<i>CostTrend</i>	-0.18	0.70	-2.06	1.53	8,185
<i>CostVol</i>	1.62	1.06	0.00	4.36	8,185

Table 3: Selected Descriptive Statistics by Strategic Cluster

Variable	Cluster 1 (22 firms)		Cluster 2 (8 firms)		Cluster 3 (13 firms)	
	Mean	Std. dev.	Mean	Std. dev.	Mean	Std. dev.
<i>Markup</i>	0.01	0.09	0.03	0.08	0.12	0.09
<i>Rank</i>	9.22	5.91	13.16	6.32	17.82	4.61
<i>SinceChange</i>	96.23	121.67	210.82	194.70	103.76	108.91
<i>NumBump</i>	1.14	3.21	2.06	4.83	1.68	3.89

Table 4: Probit Model of Timing of Price Change

Variables	Firms combined	Cluster 1	Cluster 2	Cluster 3
<i>Markup</i>	-0.020*** (0.004)	-0.016*** (0.005)	-0.018*** (0.004)	-0.007 (0.013)
<i>Markup</i> ²	0.041*** (0.013)	0.059*** (0.021)	-0.006 (0.028)	0.023 (0.048)
<i>SinceChange</i>	-0.026*** (0.004)	-0.032*** (0.007)	0.002 (0.002)	-0.016*** (0.008)
<i> NumBump </i>	0.618*** (0.123)	0.994*** (0.211)	0.226*** (0.048)	0.355*** (0.143)
<i>CostTrend</i> (when positive)	1.874** (0.873)	2.915*** (0.938)	1.785* (0.851)	-1.657 (3.317)
<i>CostTrend</i> (when negative)	-3.337*** (0.861)	-3.013** (1.360)	-0.778 (1.061)	-3.583*** (0.942)
<i>CostVol</i>	0.258 (0.216)	0.368 (0.306)	-0.511 (0.418)	1.135** (0.389)
<i>Placement</i>	0.209 (0.808)	0.209 (0.808)	-1.771*** (0.710)	1.821 (1.714)
<i>Density</i>	1.869** (0.834)	2.549** (1.098)	2.167*** (0.654)	-0.779 (1.734)
<i>PriceLow</i>	-0.032*** (0.013)	-0.047*** (0.017)	0.003 (0.011)	0.010 (0.021)
<i>PriceRange</i>	0.343*** (0.092)	0.329*** (0.129)	-0.096 (0.125)	0.498*** (0.129)
<i>AvgMarkup</i>	0.021** (0.008)	0.023* (0.012)	0.007 (0.007)	-0.026** (0.011)
<i>NewFirm</i>	-1.577** (0.631)	-1.528* (0.842)	-2.887*** (0.935)	0.409 (0.659)
Pseudo <i>R</i> ²	0.014	0.012	0.034	0.021
Observations	165,330	107,436	20,277	37,617

Notes: Dependent variable is a 0-1 indicator for price change by firm i at time t . Marginal effects reported. Heteroskedasticity-robust standard errors accounting for non-independence across observations for same firm reported in parentheses. Results for all but the markup variables (*Markup*, *Markup*², and *AvgMarkup*) scaled up by 10³ for display purposes. Statistically significant in a two-tailed test at the *10% level, **5% level, ***1% level.

Table 5: Ordinary Least Squares Model of Size of Price Change

Variables	Firms Combined			Cluster 1			Cluster 2			Cluster 3		
Constant	-0.010*** (0.003)	-0.008 (0.005)	-0.011*** (0.003)	-0.021*** (0.006)	-0.032** (0.012)	-0.068*** (0.013)	0.008 (0.011)	0.038** (0.014)				
<i>Markup</i>	-0.161*** (0.021)	-0.194*** (0.031)	-0.198*** (0.028)	-0.290*** (0.044)	-0.340* (0.148)	-0.496*** (0.133)	-0.234*** (0.038)	-0.425*** (0.085)				
<i>SinceChange</i> (rising <i>CostTrend</i>)	0.003 (0.002)	0.003 (0.002)	0.003 (0.002)	0.003 (0.002)	0.005* (0.003)	0.006* (0.003)	0.002 (0.004)	0.002 (0.004)				
<i>SinceChange</i> (falling <i>CostTrend</i>)	-0.002 (0.002)	-0.002 (0.002)	-0.003 (0.003)	-0.002 (0.003)	0.003 (0.005)	0.004 (0.004)	0.001 (0.005)	0.005 (0.004)				
<i>NumBump</i>	-0.504*** (0.060)	-0.472*** (0.060)	-0.525*** (0.081)	-0.470*** (0.081)	-0.533*** (0.118)	-0.412** (0.130)	-0.349*** (0.091)	-0.310*** (0.117)				
<i>CostTrend</i>	0.182 (0.195)	-0.044 (0.198)	0.186 (0.226)	0.223 (0.260)	0.774 (0.976)	1.962** (0.765)	-0.090 (0.397)	-0.637 (0.368)				
<i>CostVol</i>	0.272** (0.130)	0.275** (0.125)	0.157 (0.153)	0.155 (0.143)	0.743 (0.413)	0.715 (0.616)	0.296 (0.342)	0.371 (0.366)				
<i>PriceLow</i>	0.005 (0.005)	0.005 (0.005)	0.005 (0.005)	0.017** (0.006)	0.037 (0.022)	0.037 (0.022)	0.024** (0.010)	0.024** (0.010)				
<i>AvgMarkup</i>	0.073* (0.041)	0.073* (0.041)	0.183*** (0.050)	0.183*** (0.050)	0.465** (0.154)	0.465** (0.154)	0.222* (0.114)	0.222* (0.114)				
<i>NewFirm</i>	1.464*** (0.155)	1.464*** (0.155)	-1.513*** (0.225)	-1.513*** (0.225)	-0.612 (1.243)	-0.612 (1.243)	1.339*** (0.363)	1.339*** (0.363)				
R^2	0.301	0.319	0.305	0.338	0.624	0.671	0.327	0.385				
Observations	1,593	1,593	1,253	1,253	70	70	270	270				

Notes: Dependent variable is $\ln(\text{Price}_{it}) - \ln(\text{Price}_{it-1})$ for only those observations it with a non-zero price change between $t-1$ and t . Heteroskedasticity-robust standard errors accounting for non-independence across observations for same firm reported in parentheses. Results for all but the constant, *Markup*, and *AvgMarkup* scaled up by 10^2 for display purposes. Statistically significant in a two-tailed test at the *10% level, **5% level, ***1% level.

Table 6: Actual Versus Simulated Values of Average Markups

Subgroup of firms	Actual markups	Simulated markups in listed scenario			
		Maintaining actual mix	Assuming all firms in cluster 1	Assuming all firms in cluster 2	Assuming all firms in cluster 3
All firms	0.0343	0.0203			
Cluster 1 firms	0.0056	-0.0038	-0.0142		
Cluster 2 firms	0.0276	0.0441		0.1982	
Cluster 3 firms	0.1200	0.0445			0.0370

Figure 1: Example of a Pricewatch webpage


PRICE WATCH® Query = System Memory PC100 128MB

[Back To - New Components Home](#)

[Back To - "Not Exactly New" Home](#)

BRAND	PRODUCT	DESCRIPTION	PRICE	SHIP	DATE/HR	DEALER/PHONE	ST	PART#
Generic	PRICE FOR ONLINE ORDERS ONLY - 128MB PC100 SDRAM DIMM - 8ns Gold leads	- * LIMIT ONE - Easy installation - in stock	\$ 68	9.69 INSURED	10/12/00 12:40:05 AM CST	Computer Craft Inc. 800-487-4910 727-327-7559 Online Ordering	FL	MEM-128-100PCT
Generic	ONLINE ORDERS ONLY - 128MB SDRAM PC100 16x64 168pin	- * LIMIT ONE	\$ 69	INSURED\$9.95	10/11/00 10:59:56 PM CST	Connect Computers 888-277-6287 949-367-0703 Online Ordering	CA	-
Generic	PRICE FOR ONLINE ORDER - 128MB PC100 SDRAM DIMM	- * LIMIT ONE - - InStock, 16x64-Gold Leads	\$ 70	10.75	10/11/00 2:11:00 PM CST	1st Choice Memory 949-888-3810 -- P.O.'s accepted Online Ordering	CA	-
Generic	PRICE FOR ONLINE ORDER - 128mb True PC100 SDRAM EEPROM DIMM16x64 168pin 6ns/7ns/8ns Gold Leads	- * LIMIT ONE - in stock - with Lifetime Warranty	\$ 72	9.85	10/10/00 11:30:39 AM CST	pcboost.com 800-382-6678 -- P.O.'s accepted Online Ordering	CA	-
Generic	IN STOCK, 128MB PC100 3.3volt unbuffered SDRAM Gold Lead 168 Pin, 7/8ns - with Lifetime warranty	- * LIMIT ONE Not compatible with E Machine	\$ 74	10.95- UPS INSURED	10/11/00 12:44:00 PM CST	Memplus.com 877-918-6767 626-918-6767	CA	- 880060
Generic	PRICE FOR ONLINE ORDERS ONLY - 128MB True PC100 SDRAM DIMM - 8ns Gold - warranty	- * LIMIT ONE	\$ 74	10.25	10/9/00 6:53:25 PM CST	Portatech 800-487-1327 .	CA	-
House Brand	128MB PC100 3.3volt SDRAM 168 Pin, 7/8ns - with LIFETIME WARRANTY	- * LIMIT ONE	\$ 74	10.50 FedEx	10/11/00 10:20:23 AM CST	1st Compu Choice 800-345-8880 800-345-8880	OH	-
Generic	128MB 168Pin TRUE PC100 SDRAM - OEM 16X64	DIMM16x64 168pin 6ns/7ns/8ns Gold Leads	\$ 75	\$10	10/11/00 2:37:00 PM CST	Sunset Marketing, Inc. 800-397-5050 410-626-0211 -- P.O.'s accepted	MD	-
Generic	128MB 16x64 PC100 8ns SDRAM.	- * LIMIT ONE	\$ 77	10.90	10/12/00 9:37:45 AM CST	PC COST 800-877-9442 847-690-0103 Online Ordering	IL	-
Generic	IN STOCK, PC100, 128MB, 168pins DIMM NonECC, - with Lifetime warranty	- * LIMIT 5	\$ 77	\$10.95 & UP For UPS Ground	10/9/00 5:11:10 PM CST	Augustus Technology, Inc 877-468-5181 909-468-1883 Online Ordering	CA	-
Generic	128MB PC100 8NS 16x64 SDRAM - one year warranty	- * LIMIT ONE	\$ 78	Ups Ground \$10.62	10/11/00 5:16:36 PM CST	Computer Super Sale 800-305-4930 847-640-9710 Online Ordering	IL	-
Generic	PRICE FOR ONLINE ORDERS ONLY - PC100 128MB NonBuffered, NonECC 16x64 SDRAM DIMM 3.3V 8ns mem module	- * LIMIT ONE - with lifetime warranty	\$ 78	10.95	10/5/00 6:29:59 PM CST	Jazz Technology USA, LLC 888-485-8872 909-869-8859	CA	ME-GBP100128

Figure 2: Effect of Rank on Demand



Notes: Estimates of effect of rank on demand from Ellison and Ellison (2009) using two different specifications. Solid line is from a regression of the natural log of quantity on linear rank. Dots are coefficients from a more flexible specification with rank fixed effects; vertical bars provide 95% confidence interval around fixed effects.

Figure 3: Markup histogram

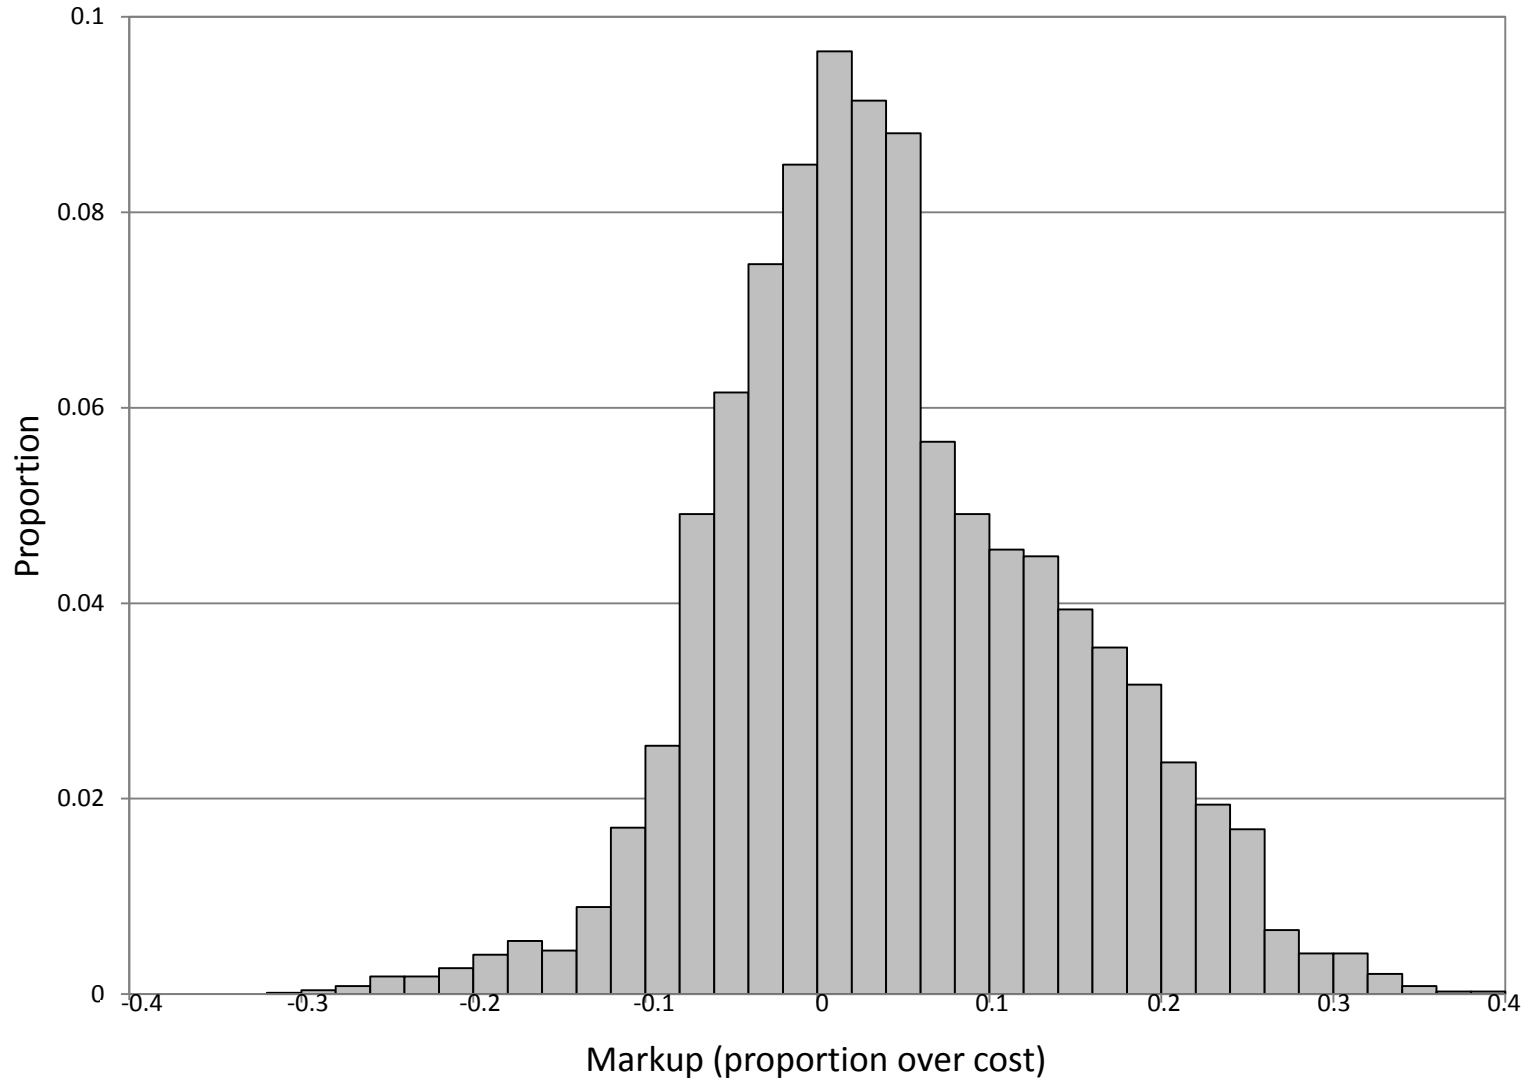


Figure 4: Markups over time




Figure 5: Price series for representative firms and months

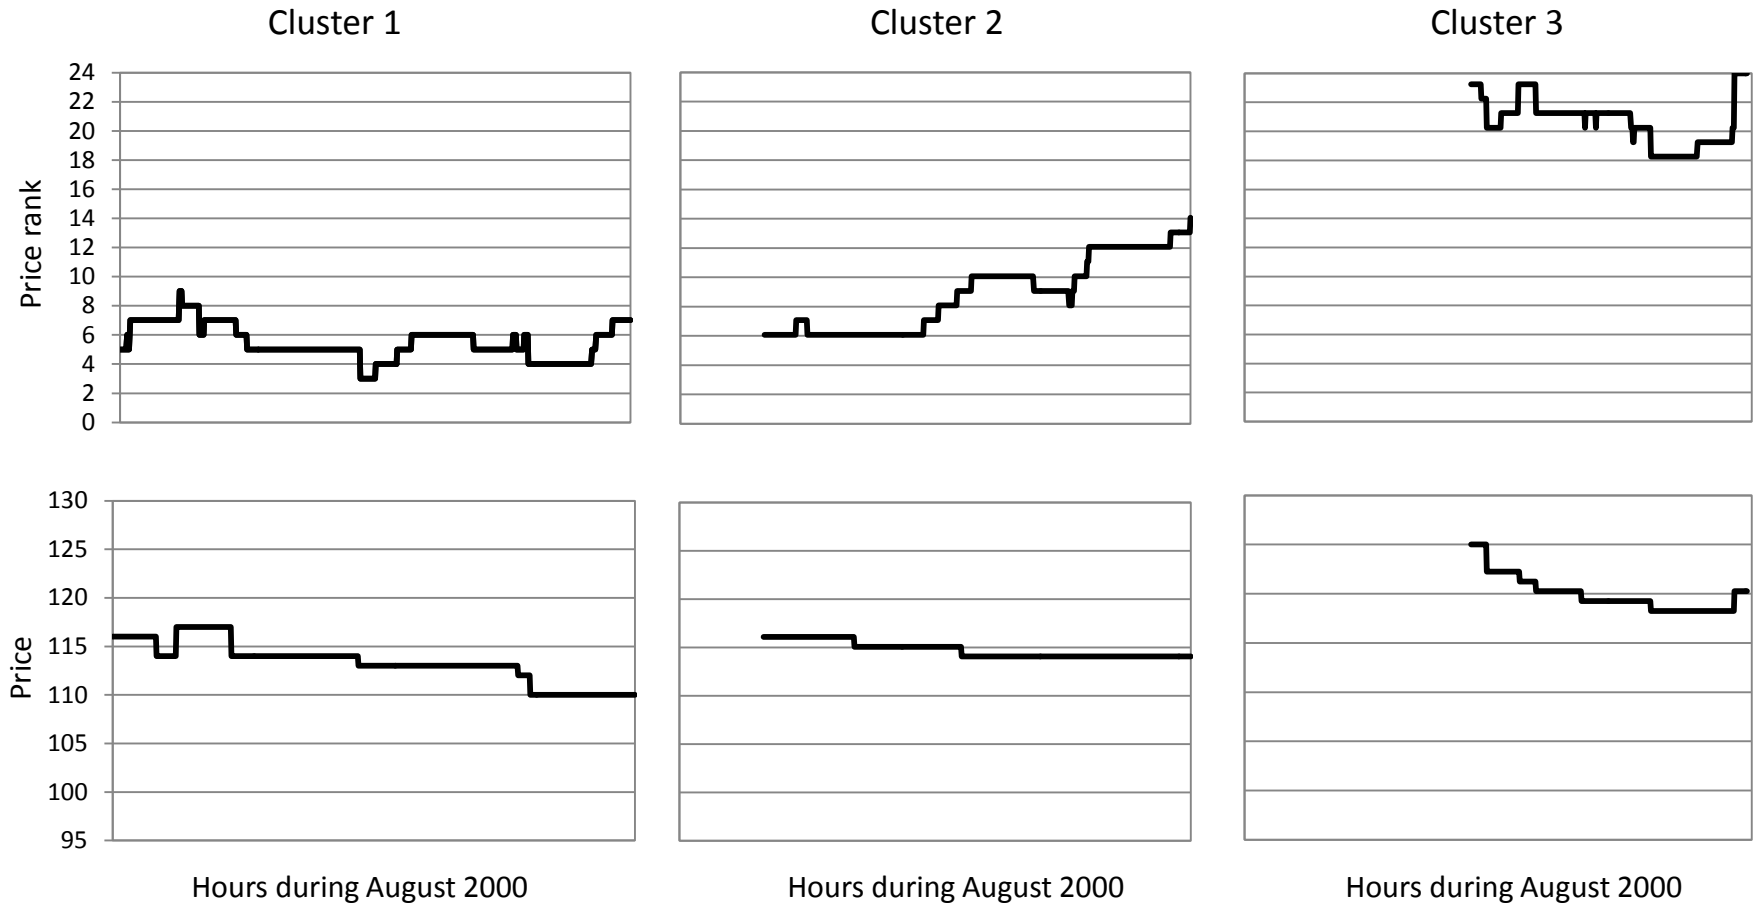
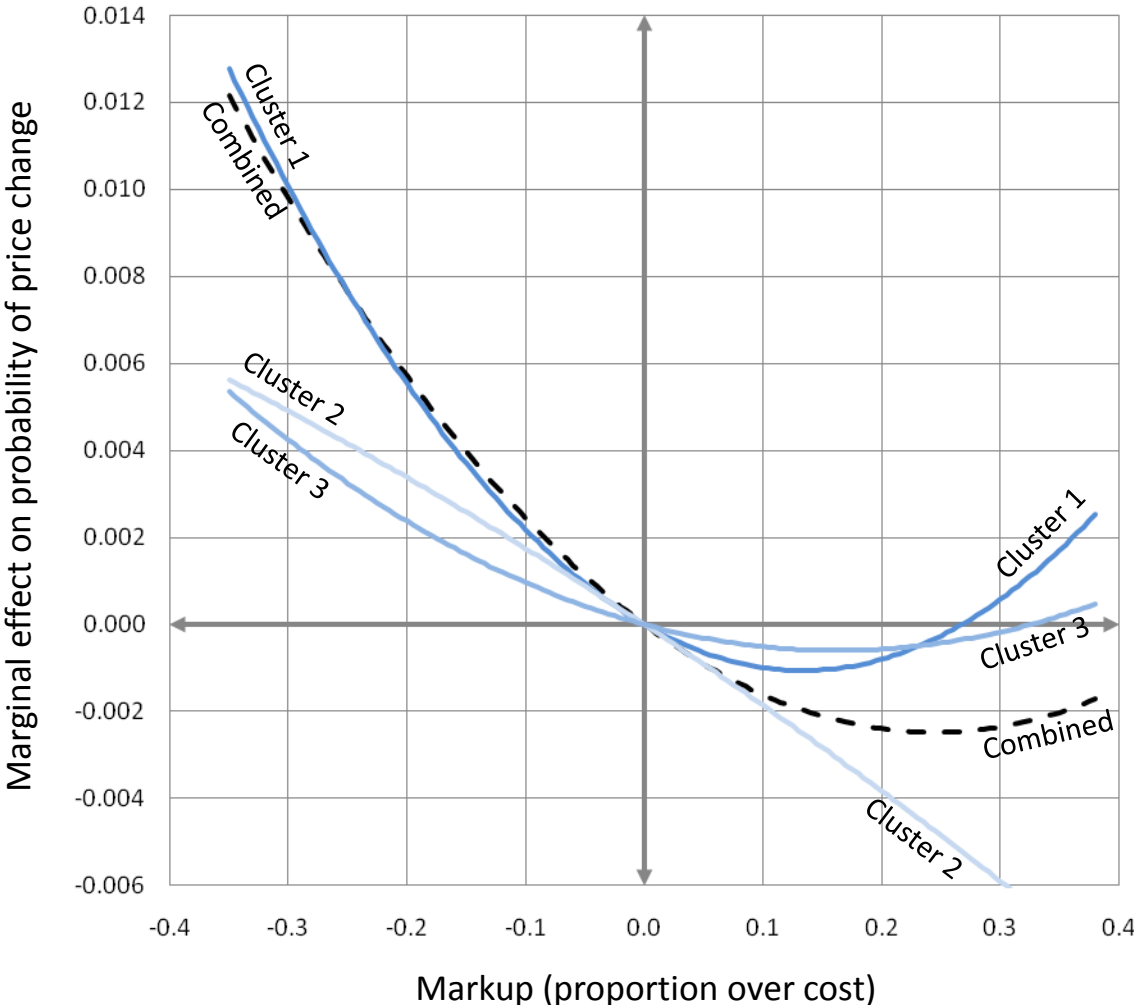


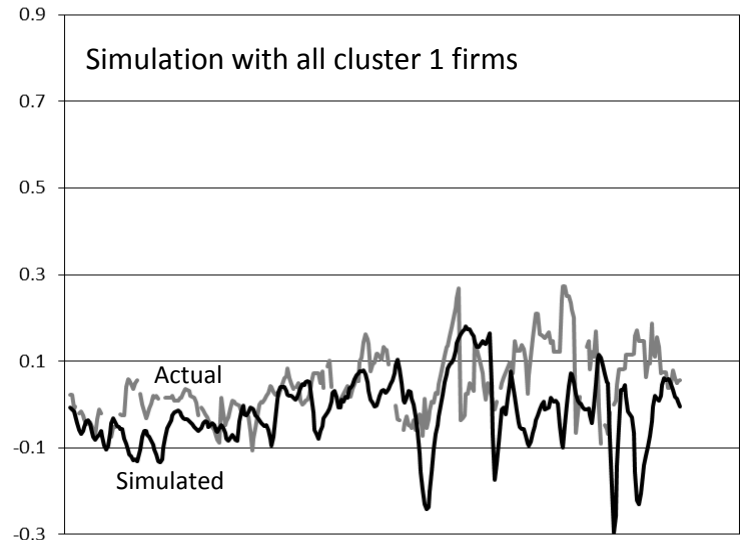
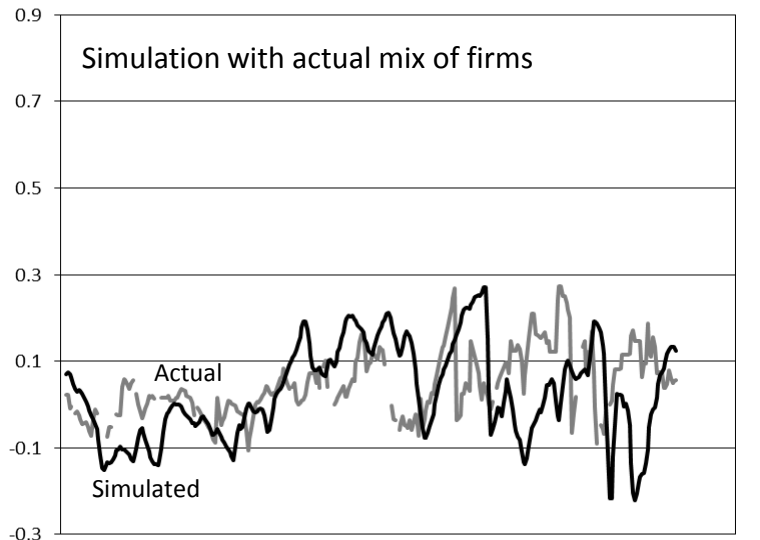


Figure 6: Quadratic marginal effect of markup on probability of price change



Note: Based on linear and squared markup coefficients from parsimonious models in Table 4 probits.

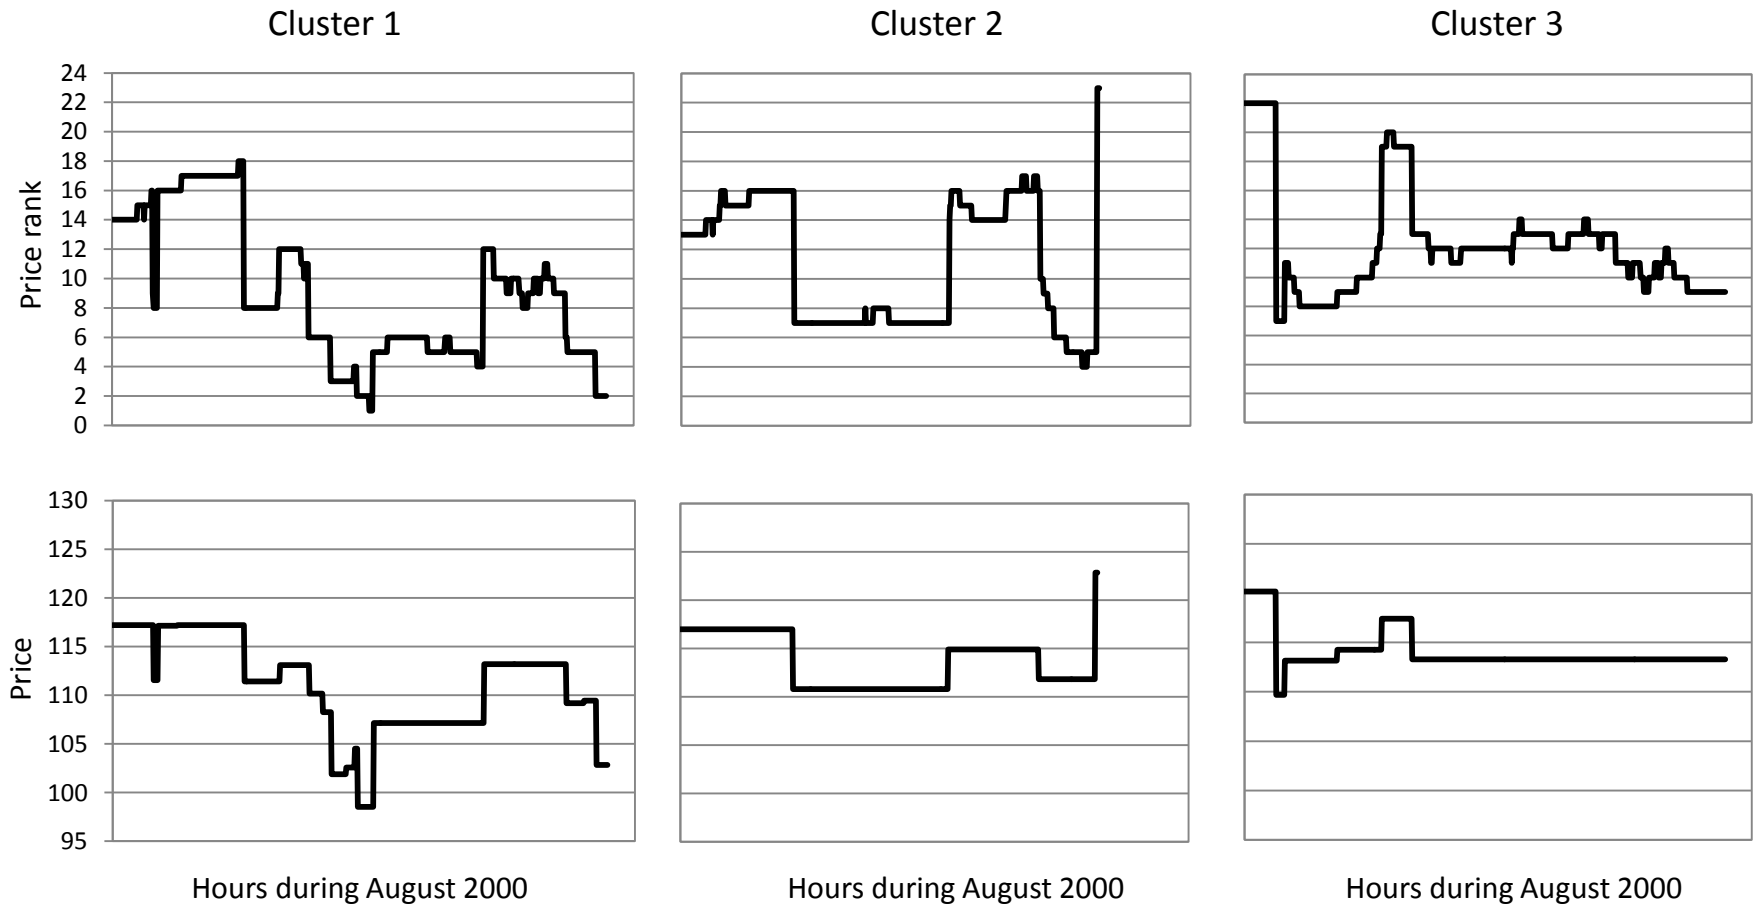
Figure 7: Simulated Versus Actual Markups for the Median Firm Over Time



Hour

Hour

Figure 8: Price series for representative simulated firm and month



Note: Simulation maintains actual mix of clusters.